

**Questioning Code:
Courage in the Age of Artificial Intelligence
Alondra Nelson**

Amherst College Commencement
Amherst, MA
May 24, 2025

As Prepared for Delivery

President Elliott, faculty, family, fellow honorands, dear graduates:

This weekend, we celebrate not just the culmination of years of study, but the beginning of a new journey – and you, graduates, as shapers of our collective future.

What brought you here is courage.

Courage.

Courage is not the absence of fear. It is action in the *face* of fear. It is the willingness to stand firm when the ground beneath you trembles. And the determination to speak truth even when there is more safety in silence.

Courage has worn many faces throughout history.

It is Galileo Galilei insisting that the Earth revolved around the sun – despite the mortal threats of the Inquisition. It is Rosa Parks, who refused to surrender her seat, risking arrest, and catalyzing a movement that would transform a nation. It is the piercing voice of Billie Holiday, years before the civil rights era, singing the words of Lewis Allen, telling of the *strange fruit* – the bodies of Black lynching victims – hanging from trees across the South.

Courage is in our art – the lifeblood of centuries of human creativity, the object of our novels and music, and our plays. It is Antigone, defying King Creon to honor her brother, choosing moral law over state decree. Courage is in those long-ago memories, and those most essential of stories.

But it is also right here. I mean really, right here, among people who have walked the grounds of Amherst College.

When I was writing my first book, *Body and Soul*, about independent and audacious responses to unequal healthcare access and medical discrimination in the 1960s and 1970s, I contextualized these efforts with the work of an extraordinary historical actor, Dr. Charles Drew, Amherst Class of 1926. (And one of only 13 African Americans in the College's student population at the time). Dr. Drew was a pioneering physician-researcher, who revolutionized blood preservation techniques that saved countless lives during World War II.

Dr. Drew found himself in an impossible position: the very blood banking programs he had created to save lives – including at the American Red Cross– required that blood donations from Black donors be kept separate from white donors, despite having no scientific basis.

Rather than stay silent to protect his career, Drew spoke out publicly against these segregation policies, even though criticizing the organizations he was helping could have ended his professional opportunities. His willingness to risk everything to challenge institutional racism while continuing his life-saving medical innovations showed what it means to live by your principles, even when the cost is high.

We occasionally witness these traits in our own lives – in the student who asks the question no one else dares to voice, in the researcher who pursues an unpopular hypothesis, in the friend who speaks a difficult truth with love when it would be easier to remain silent.

But it is rare. It is hard. And often it is recognized by society, and honored as a virtue, only after the moment of courage has passed. The formula here – employed by each of these figures – is to feel the currents of one’s life moving in one direction and muster the strength to walk in another direction.

The currents of this moment are fast and furious. You are graduating into a world of intersecting crises: climate change, biodiversity loss, political polarization, economic inequality, and persistent discrimination. Not to mention rising authoritarianism, in this country and around the world.

On the one hand, these problems are not altogether different from the world into which your parents and grandparents came of age. There is always war of some kind raging in our world. And there is always upheaval, economic and social. Indeed, there is always inequality – and people dedicated to fighting it.

But there is at least one factor that is different. One current in particular that could exacerbate these crises, or help address them. One force moving rapidly, sweeping all of us up with it: Artificial intelligence.

You are graduating into the age of algorithmic governance and automated decision-making. Of new, powerful technological tools. This is what I’d like to focus on this morning: Courage in the time of artificial intelligence.

What does courage look like in an age of unlimited choices – when you can ask AI to write your essay, solve your problem set, or craft your job application? Just because you can use these tools, does that mean you should? In an age of such technological power, how do we hold onto our own moral compass and decide not just what we can do, but what we *should* do?

Before we can make wise choices about these tools, we need to understand what they really are and how they came to have such power over our lives. So allow me, if you don’t mind, to step into professor mode for just a moment with a brief history lesson.

Artificial intelligence began in the 1950s when computer scientists first dreamed of creating machines that could handle complex tasks like pattern recognition, problem solving, and data

analysis – leading to early breakthroughs in programs that could play games and make automated decisions in fields like medicine and finance.

After decades of ups and downs – including periods when AI research nearly disappeared because it promised too much and delivered too little – the field exploded in the 2010s thanks to massive amounts of data, powerful computing, and breakthroughs in machine learning that let computers identify and classify images, convert speech to text, and even beat world champions at complex games like Go.

Today, AI has evolved from a research curiosity into the technology powering your smartphone's voice assistant, Netflix recommendations, and large language models like ChatGPT, Claude, and Llama that can generate text, images, and code with remarkable sophistication and coherence.

But artificial intelligence has also evolved from a theoretical concept to a ubiquitous force – a *powerful* force – reshaping how we work, create, communicate, and even how we understand ourselves.

AI systems now create art, recommend medical treatments, determine who gets loans, and increasingly influence who gets hired, promoted, or even admitted to institutions like this one. Large language models can analyze vast datasets and deliver conclusions faster than human researchers. Recommendation algorithms curate our information environment, determining what news we see, what ideas we encounter, and increasingly, what opinions we form.

And *get this*: Some people even use AI to *generate essays*. But I know nobody at Amherst College would know anything about that.

But beyond the classroom, these systems are making far more consequential decisions. In my work leading the White House effort to draft the Blueprint for an AI Bill of Rights, we heard from countless communities concerned about AI systems making decisions about their health, their jobs, even their liberty – without accountability or recourse.

These concerns are real, and they demand our attention and action.

The technology companies developing these systems have made big promises: Unprecedented efficiency, productivity, and progress. They tell us that AI will solve humanity's greatest challenges – climate change, disease, educational inequality. They assure us that algorithmic decision-making will rid the world of human bias and bring opportunity to everyone.

As such, they suggest that there is only one rational path forward: to embrace automation in all its forms, to follow the code, wherever it may lead.

Perhaps.

But I'm not so sure. And I would invite you to nurture your own skepticism, as well.

It is true that AI tools have demonstrated remarkable capabilities that might benefit humanity. But there are also dangerous fallacies emerging.

Perhaps you've heard that technological progress is inevitable, that human judgment is obsolete. To which I would ask you: Have you heard of the Windows phone? Have you used the Apple Lisa computer?

When Apple released the Lisa computer in 1983, it was the first personal computer with a graphical user interface and mouse – technology design that it directly inspired the Macintosh and eventually every computer you've ever used. Yet it failed because people weren't ready to embrace such an expensive, unfamiliar way of computing.

Similarly, Microsoft's Windows Phone featured a superior user interface with innovative live tiles and better performance than early Android phones, but it couldn't overcome the fact that people had already formed habits and communities around iOS and Android, and simply didn't want to learn a new system.

These failures show that the best technology isn't automatically adopted. Instead, its diffusion depends on whether people choose to adopt it, whether it fits into their existing social networks and daily routines, and whether society is ready for the change it represents. Understanding this reality is crucial because it means we have agency in shaping how technology develops and gets used, rather than simply being passive recipients of whatever engineers and companies decide to build.

Perhaps you've heard that code is neutral.

Code, the software that structures our digital experiences, actually functions as a powerful regulator that shapes fundamental values like privacy, free speech, and individual liberties. Unlike neutral tools, code reflects specific choices about what behaviors to enable or restrict, such as whether identification systems protect anonymity or reveal personal information.

Consider how Facebook's ad targeting algorithms automatically excluded women from seeing STEM job advertisements and prevented older workers from seeing certain employment opportunities. Or consider the facial recognition technology that wrongly identified Robert Williams, an innocent Michigan father and husband, as having committed a crime – resulting in his humiliation, jail time, and financial burden.

When code is harmful, we should question it.

Programmers, developers, and corporations are effectively writing the rules that govern our digital lives. The belief in code's neutrality blinds us to how these architectural decisions embed particular values and interests, allowing private coders to determine our freedoms, often without democratic oversight or constitutional protections.

Certainly you've absorbed by now the narrative – from our media, from our markets, and from our country's business executives – that efficiency is *always* the highest value. And that questioning the trajectory of technological development is foolish or futile.

I don't buy it. Any of it. Neither do many of my friends and colleagues who study technology, its history, and its impact on humanity.

And so today I invite you to do something that requires courage. As you step out into the age of artificial intelligence, I challenge you to question the code.

I want you to interrogate algorithms, to challenge the narrative that technology's current path is predetermined.

I want you not just to ask *can* we do this ... but *should* we do it? Not just "How fast?" but "In what direction?" Not just "Who benefits?" but "Who gets to decide?"

What it requires in order to press against these powerful currents is to get clear on one's own values. To know and to articulate a different code, your personal code of conduct.

Each of you has developed the most human of capacities: critical thinking. At Amherst, you've learned to analyze arguments, identify assumptions, evaluate evidence, and construct reasoned perspectives. These skills are the foundation of intellectual freedom. They are what allow you to form independent judgments rather than simply absorbing the judgments of others.

Now, we are told that artificial intelligence can "think" for us. The convenience is seductive. Why struggle with a difficult text when AI can summarize it? Why wrestle with a complex problem when an algorithm can propose solutions? Why navigate the messy landscape of diverse viewpoints when personalized feeds can confirm what we already believe?

I don't want to trivialize this impulse to just use AI. Independent thought is the harder path. When we outsource our thinking to algorithms, we're surrendering our intellectual autonomy to systems that shape not just our answers, but how we think, what questions we ask, and ultimately who we become as thinkers. Independent thought requires you to do the mental heavy lifting – grappling with uncertainty, spending time researching and analyzing, and risking being wrong.

When algorithmic shortcuts are available everywhere, it takes integrity to say, "I will read this myself, think this through myself, reach my own conclusions." It will require frequent personal decisions to engage directly with primary sources when confronted with the option of machine-generated summaries that may flatten nuance or miss crucial context.

Critical thinking in the age of AI means understanding not just how and when to use AI tools, but how they work, what biases they may encode, and what limitations they inherently possess. It means building a practice of asking questions: What data was this system trained on? Whose experiences and perspectives might be overrepresented or underrepresented in that training data? What assumptions are built into the algorithm itself? Who benefits from its deployment?

It means recognizing that large language models are essentially sophisticated pattern-matching systems trained on human-generated texts—they do not "understand" or "reason" in the way humans do.

And so to question code – as I have proposed – is not anti-technology. It is pro-humanity. It acknowledges that technology is not neutral but reflects the values, priorities, and power structures of the society that creates it.

It takes courage to insist that the progress of AI serve human goals and freedoms. It takes courage to defend what makes us human in an age increasingly defined by optimization. And it takes courage to defend what makes us human in an age increasingly defined by optimization.

The logic of the age of AI is fundamentally a logic of optimization – identifying patterns, maximizing efficiency, reducing friction, eliminating unpredictability. This approach may be valuable in some contexts. But human life is not merely a series of optimization problems. Sometimes I wish it were. That would be much easier.

And yet, our greatest achievements, our deepest connections, and our most profound insights almost never emerge from optimization. They come from exploration. They come from experimentation, from trying and failing and getting things wrong and wandering down new paths, and learning.

Consider the origins of the World Wide Web. Tim Berners-Lee was not trying to optimize information sharing when he invented the precursor to today's internet. He was frustrated by the difficulty of sharing research at the European Organization for Nuclear Research or CERN and set out to create an open, decentralized space for researchers. An optimization approach would have built in more security, control mechanisms, and profit models. Instead, his "inefficient" design and giving it away for free created the foundation for the entire modern internet.

What makes us human is not just our capacity for rational thought but our capacity for wonder, for ethical judgment, for aesthetic appreciation, for empathy – qualities that resist algorithmic encoding. To preserve space for uncertainty, ambiguity, and exploration when technology pushes toward premature certainty and closure. That, too, takes courage.

Consider what happens when you encounter injustice – maybe witnessing someone being treated unfairly because of their race, religion, or nationality. Your ethical response isn't calculated through data points or probability matrices. It springs from your moral imagination, your ability to feel another's pain as your own, and your capacity to envision a more just world. These ethical judgments require the uniquely human ability to see beyond patterns to possibilities.

Or, think about a painting, a song, a poem, a play that stopped you in your tracks. The algorithm that recommended it might predict you'd "like" it based on your history, but it cannot experience the way that particular combination of colors or words or tones opens something new in your soul, or connects you to a truth you couldn't articulate before.

While questioning code and defending humanity, we must simultaneously embrace a third form of courage: The courage to learn, understand, and shape these new technologies rather than simply accepting or rejecting them wholesale.

AI is not going away. These tools will continue to evolve and integrate into our personal and professional lives. Those who will thrive in this environment will be neither uncritical adopters nor reflexive resisters, but principled navigators – people who understand both the capabilities and limitations of AI systems and can use them as tools for human purposes without being used by them.

Learning new technologies takes bravery. For those of us who have demonstrated mastery in our chosen fields, this return to novice status can be uncomfortable. Mastering emerging tools asks us to risk the vulnerability of being beginners again.

You don't need to become a computer scientist, though we certainly need more ethically-minded ones. But all of us must develop enough technical understanding to ask informed questions: What is this technology designed to do? How does it work? Where might it fail? What human activities does it support, and what human capacities might it diminish?

This technical acuity must extend beyond technical understanding to include historical and social context. The history of technology is full of innovations that promised benefits while obscuring their costs.

Consider the One Laptop Per Child campaign. It promised to revolutionize education in developing countries by giving every child a computer, claiming it would bridge the digital divide, democratize knowledge, and lift communities out of poverty. But the hidden costs were considerable: funding was diverted from basic educational needs like teachers, textbooks, and infrastructure. Worse, it fostered the illusion that technology alone could solve deeply rooted social and economic challenges, often ignoring local languages, contexts, and priorities.

Or take social media. It was marketed as a tool to connect the world and democratize information. But its concealed costs have included negative mental health impacts on young people, the rapid spread of misinformation, and the erosion of personal privacy through constant surveillance of our digital footprints.

These examples remind us to reject technological determinism, the belief that technology follows a fixed, inevitable path beyond human influence. Technologies are not neutral forces; they are human creations, shaped by human decisions, priorities, regulations, and resistance.

They can be designed differently, deployed differently, regulated differently, or in some cases, not deployed at all if their social costs outweigh their benefits.

As you enter or return to the workforce, remember that you are in the driver's seat. You do not have to be passive recipients of AI tools but can be active shapers of how they are integrated into your organizations, communities, and lives. Advocate for responsible AI practices, for meaningful human oversight, for ethical guidelines that prioritize human well-being over narrow efficiency metrics.

Question what constitutes meaningful work when many traditional tasks can be automated. Ensure that the economic benefits of AI do not flow exclusively to those who own the technology while leaving others behind. And as you rise in stature and in power, make sure that you are building a future where technology serves justice, sustainability, and human wellbeing.

Building a just future with AI requires confronting how these technologies can amplify existing inequalities. Facial recognition systems that work less accurately for darker-skinned faces. Hiring algorithms that reproduce gender bias. Risk assessment tools that perpetuate racial disparities in criminal justice. Surveillance technologies that disproportionately target marginalized

communities. These are not merely technical problems but social and political ones that require courage to address.

When we question code, we must also question its consequences, not just the social and political impacts I've described, but the physical ones as well. Every algorithm runs on hardware. Every AI query requires electricity. Every data center needs cooling. Building a livable future with AI means recognizing that questioning code isn't just about bias and fairness – it's also about confronting the hidden environmental costs of these technologies. It takes enormous amounts of energy and water to train large AI models. The extraction of rare earth minerals for computing hardware creates environmental devastation in communities rarely consulted about these costs.

Question code when it is extractive: The environmental cost of AI is staggering. In just one year, the power needed to run AI systems nearly doubled. If all the world's data centers were a single country, they would use more electricity than most nations—more than Saudi Arabia, and by 2026, more than Germany. Every time you ask a chatbot question or generate an AI image, you're drawing power from the same electrical grid that lights your home.

The water consumption required to cool AI infrastructure represents another hidden environmental burden. Training, deploying, and fine-tuning generative AI models generates enormous heat, demanding massive amounts of water for cooling systems that can deplete municipal water reserves and destabilize local ecosystems.

Right now, as I speak, there's a community in Southwest Memphis, Boxtown, founded by formerly enslaved people, where residents are watching the wealthiest man in the world burn through their air and water to train his AI systems. Elon Musk's xAI facility runs 35 methane gas turbines, making it likely the largest industrial polluter in Memphis, in a community where cancer rates already exceed the national average by four times.

When you generate an AI image or ask a chatbot a question, remember that somewhere, a real community – often a community of color, often a community already bearing the burden of environmental pollution – may be paying the price with their lungs, their water, their children's health. Every algorithm has an address. Every data center has neighbors. And those neighbors rarely get to vote on whether the convenience of our AI queries is worth the cost of their health.

Yet AI also offers tools for environmental monitoring, renewable energy optimization, and more efficient resource use. Recent breakthroughs like DeepSeek's reasoning models suggest we can develop AI systems that require dramatically fewer resources, proving that the current trajectory of ever-increasing consumption is not inevitable but a choice. The question is not whether we use these technologies, but whether we have the will to direct them toward sustainability rather than extraction, toward sufficiency rather than endless growth.

Some of you will contribute to this building directly. You'll be the computer scientists developing more transparent algorithms, designing more inclusive datasets, creating governance structures that ensure the use of AI benefits humanity broadly. Others will contribute through advocacy, education, art, policy, or simply through the conscious choices you make as citizens and consumers.

But sitting this one out will not be an option. As we step out into a future inflected by artificial intelligence, we will all have choices to make.

So how do we actually *make* these choices? How do we have the courage to ask hard questions and hold onto principle when the currents of history – indeed, the currents of this technology – are moving so rapidly and so decisively?

I believe there are five steps we can take.

First, we practice intellectual courage – by creating space for deep thinking. This means occasionally disconnecting from the constant flow of information to reflect, synthesize, and form independent judgments. It means reading primary sources rather than just summaries, engaging with viewpoints that challenge your assumptions, and being willing to change your mind when evidence warrants.

Second, we practice social courage – by building community. Questioning powerful technologies is easier to sustain when you're not alone, when you're not scrolling alone through feeds to amplify emotion, but instead engaged in real conversations with real people. Find or create spaces—both digital and physical—where people can share concerns, develop alternatives, and support each other in ethical technology practices.

Third, we practice political courage – by governing our technology as we govern other aspects of our shared life. Technologies are shaped not just by their creators but by regulations, standards, and public opinion. Vote for candidates who understand these issues. Participate in public comment periods for proposed regulations. Support organizations advocating for responsible uses of AI.

Fourth, we practice personal courage – through being mindful about our technology choices. Which AI tools do you use? For what purposes? With what limitations? When do you choose human connection over technological convenience? These seemingly small decisions, multiplied across millions of people, shape our technological future as surely as any corporate strategy or government policy.

Finally, we practice visionary courage – by imagining alternatives. Our current technological paradigm can seem inevitable, permanent, unchangeable. But it is not. The same imagination that allows us to envision dystopian futures also empowers us to dream of better ones: technologies that enhance rather than replace human connection, AI systems that extend rather than diminish human creativity, digital spaces that foster democracy rather than division. The courage to imagine different relationships between humanity and technology, and to work toward making those visions real, may be the most important courage of all.

Envision the year 2035: Will we live in a world where uses of AI help communities thrive, or where communities like Boxtown sacrifice their children's lungs for Silicon Valley's profits?

You are entering a world that both needs and challenges your courage. You will face pressures to conform, to accept technological changes without question, to prioritize career advancement over ethical concerns. There will be moments when speaking up seems too risky, when asking hard

questions feels futile, when the systems you hope to change seem immovable. We have all been there.

In those moments, I would ask you to remember that courage is not a quality you either have or lack – it is a practice you develop through repeated choices. Each time you ask a difficult question, raise an ethical concern, learn a new skill, or take a stand for human dignity, you strengthen your capacity for courage. Each time you connect with others, you create conditions where collective courage becomes possible.

It does not require certainty. You don't need to have all the answers about how AI should be developed or governed. You simply need the willingness to ask questions, to learn from mistakes, and to keep human rights and opportunities at the center of your technological decisions.

As you navigate this age of artificial intelligence, I hope you will be neither technophobic nor techno-utopian, but technologically wise, principled, and courageous – understanding these tools deeply enough to use them for human purposes without being used by them. I hope you will neither despair at technology's challenges nor abdicate responsibility to its creators, but engage as active participants in shaping our technological future.

I hope you will question the code.

Lastly, “code” has another meaning. It is a set of principles that helps shape our behavior. And so questioning code also means examining the values and ethics that guide our individual and collective choices about technology.

What code will you live by in this age of artificial intelligence? Will you prioritize convenience over autonomy, efficiency over justice, acceleration over sustainability? Or will you have the courage to insist that technology is better?

The answers to these questions will not come from AI systems, no matter how sophisticated. They will come from you—from your critical thinking, your moral imagination, your collaborative action, and your courage to question, defend, learn, and build.

Question code. Defend humanity. Learn continuously. Build justly. And in doing so, help create a future where technology serves human wellbeing rather than diminishing it.

Thank you for listening.